AI-ASSISTED EXPLOIT DEVELOPMENT

USING LARGE LANGUAGE MODELS TO ACCELERATE VULNERABILITY RESEARCH, REVERSE ENGINEERING, AND OFFENSIVE SECURITY

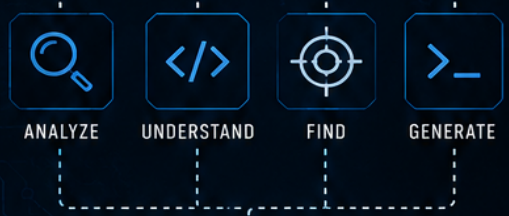

EXPLOIT

IDA PRO	Disassemble and analyze binaries with AI
GHIDRA	Reverse engineer smarter, faster
BINARY NINJA	Navigate and understand complex code
BURP SUITE	Find and exploit web vulnerabilities
DEBUGGERS	Automate analysis and runtime inspection
FUZZERS	AI-guided fuzzing and crash triage
MCP & TOOLING	Integrate AI into your offensive pipelines

10110
01001
BINARY ANALYSIS

WEB APP TESTING

MALWARE RESEARCH

VULNERABILITY DISCOVERY

OPERATIONAL SECURITY

STEVE T.

AI-Assisted Exploit Development

Using Large Language Models to Accelerate Vulnerability Research, Reverse Engineering, and Offensive Security

Steve T. Team Publications

This book is available at https://leanpub.com/ai-assistedexploitdevelopment

This version was published on 2026-07-03

Contents

Using Large Language Models to Accelerate Vulnerability Research, Reverse Engineering, and Offensive Security

About this book: This book is a practical guide for security professionals who want to integrate large language models and AI-powered tools into their vulnerability research workflows. It covers how to use AI with platforms like IDA Pro, Ghidra, Binary Ninja, Burp Suite, debuggers, fuzzers, and MCP-enabled tooling to analyze code, automate repetitive tasks, identify vulnerabilities, generate proof-of-concept exploits, and streamline exploit development pipelines. Through real-world case studies, hands-on examples, and offensive security exercises, readers learn how to combine human expertise with AI assistance to improve efficiency across binary analysis, web application testing, malware research, and vulnerability discovery, while understanding the limitations, validation requirements, and operational security considerations of AI-generated results. Designed for exploit developers, penetration testers, reverse engineers, red team operators, and security researchers.

Dedication

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

Introduction: The New Exploit Developer's Toolkit

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

How This Book Is Organized

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

Who This Book Is For

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

What This Book Does Not Cover

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

A Note on Ethics and Responsibility

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

How to Use This Book

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

Chapter 1: The New Exploit Developer's Toolkit

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

The New Reality: What Has Changed

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

A Second Case: Finding a Linux Kernel Zero-Day with o3

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

A Third Case: Autonomous Exploit Generation Against Modern Mitigations

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

What AI Does Well: The Force Multiplier Effect

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

Thinking About AI as a Tool, Not a Replacement

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

A Note on Ethics and Responsibility

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

Chapter 2: Foundations of AI-Assisted Security Research

How Transformers Work: Enough to Understand What They Can and Cannot Do

The Context Window and Its Security Implications

Probabilistic Generation and Nondeterminism

Prompt-Based Inference vs. Fine-Tuning: Choosing Your Approach

Prompt Engineering Paradigms for Security

Evidence-First Prompting

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

Chain-of-Thought Prompting

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

Structured Output Prompting

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

Constraint-Based Prompting

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

The Trust Boundary: Formalizing Human-AI Interaction

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

Key Principles for Working with AI in Security Workflows

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

Hands-On Lab: Evaluating LLM Output Reliability

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

Chapter 3: The Evolution of Automation in Security Research

The Early Era: Signature-Based Detection and Rule Matching

The Structural Analysis Era: Static and Dynamic Program Analysis

The Generative AI Era: LLMs as Security Co-Pilots

The Quantitative Leap: What AI Adds to Each Era

Why AI Represents a Qualitative Shift

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

The Limits of Automation: What No Paradigm Can Do

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

Hands-On Lab: Tracing the Evolution of Automation

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

Chapter 4: Setting Up an AI-Augmented Research Environment

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

The Landscape of AI Tools for Security Research

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

Local LLM Deployment Options

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

Model Selection for Security Tasks

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

Cloud API Considerations

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

Integrating AI with Reverse Engineering Platforms

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

The Model Context Protocol: Connecting AI to Tools

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

Operational Security for AI-Assisted Workflows

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

Hands-On Lab: Deploying a Local AI Research Environment

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

Chapter 5: AI in Binary Analysis and Reverse Engineering

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

Automated Function Identification and Naming

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

Decompilation Refinement with LLMs

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

Guided Reverse Engineering: Asking the Right Questions

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

Case Study: AI-Assisted Analysis of XLoader 8.0

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

Limitations and Failure Modes

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

Hands-On Lab: AI-Assisted Binary Analysis with Evidence-First Prompting

Chapter 6: Vulnerability Discovery with AI and Fuzzing

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

The State of AI-Guided Fuzzing

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

Key Research: FuzzGPT, CHATAFL, TitanFuzz, and Beyond

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

Performance Comparison: Traditional vs. AI-Guided Fuzzing

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

The Validation Problem

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

Static Analysis Augmentation with AI

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

Web Application Vulnerability Discovery

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

Case Study: Sean Heelan's CVE-2025-37899 Discovery

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

Hands-On Lab: AI-Guided Protocol Fuzzing with CHATAFL Principles

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

Chapter 7: Proof-of-Concept Exploit Development

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

The Landscape: AI-Powered Exploit Development Platforms

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

Shellcode Generation

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

ROP Chain Construction

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

Mitigation Bypass Techniques

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

Case Study: Sean Heelan's QuickJS Exploit Experiments

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

Safety, Ethics, and Responsible Disclosure

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

Hands-On Lab: Building an AI-Assisted Shellcode Test Harness

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

Chapter 8: Malware Analysis and AI-Assisted Triage

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

The Malware Analysis Pipeline

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

AI-Powered Initial Triage

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

Benchmarking LLMs for Malware Triage

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

Automated YARA Rule Generation

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

Unpacking and Deobfuscation with AI

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

Case Study: Rapid Triage of a New Malware Family

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

Limitations and Risks

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

Hands-On Lab: AI-Assisted Malware Triage with YARA Generation

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

Chapter 9: Web Application Exploit Development with AI

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

The Web Application Attack Surface

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

Burp Suite and AI Integration

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

AI-Assisted Vulnerability Discovery Workflows

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

Common Pitfalls and How to Avoid Them

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

Hands-On Lab: AI-Assisted Web Application Vulnerability Discovery

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

Chapter 10: Advanced Workflows and Automation

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

Building AI-Powered RE/ED Pipelines

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

MCP-Enabled Tooling: Connecting LLMs to Debuggers and Disassemblers

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

Automated Report Generation and Documentation

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

Multi-Agent Workflows: Orchestrating Multiple AI Models

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

Case Study: End-to-End AI-Assisted Vulnerability Research Pipeline

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

Scaling Your Workflow: From Individual Researcher to Team Operations

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

Chapter 11: Operational Security and Risk Management

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

The Data Flow Problem: What Goes Where, When, and Why

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

Practitioner Pushback: Why Many Reverse Engineers Refuse AI

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

Operational Security Protocols: A Practical Framework

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

The Economic Reality of AI-Assisted Security Work

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

Legal and Ethical Considerations

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

Hands-On Lab: Building an OpSec-Secure AI Research Environment

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

Chapter 12: The Future of AI-Assisted Exploit Development

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

Multimodal Models and Vision-Based Reverse Engineering

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

Specialized Security LLMs: The Domain-Specific Revolution

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

The Arms Race: AI-Assisted Attacks vs. AI-Assisted Defense

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

Skills That Will Remain Valuable as AI Matures

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

What the Next Five Years Look Like

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

Preparing for the Future: A Strategic Framework

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

Chapter 13: Becoming the AI-Augmented Expert

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

The Expert's Mental Model: A Decision Framework

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

The Iterative Refinement Loop

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

Building Your Personal AI Toolkit

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

The Professional's Code: Ethical Guidelines for AI-Assisted Security Research

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

The Long Game: Sustaining Effectiveness Over Time

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

Chapter 14: Conclusion

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

References

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

Appendix A: Quick Reference Card — Model Selection Guide

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

Appendix B: Prompt Templates — Quick Reference

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

Appendix C: Common Failure Modes — Quick Reference

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.

Appendix D: Glossary of Key Terms

This content is not available in the sample book. The book can be purchased on Leanpub at https://leanpub.com/ai-assistedexploitdevelopment.